

SLOW and ERROR: from Oslo to DC

**draft-ietf-pilc-slow-02.txt and
draft-ietf-pilc-error-02.txt**

Changes made to both drafts:

Two updates to each draft since Oslo:

Version 01 - reflected comments from Oslo

Version 02 - reflected discussions since Oslo

Both titles have changed

Both abstracts have changed

Background now less history, more technology

Lots of editorial cleanup and minor corrections

SLOW: Significant Changes in 01

Recommended limited receiver window

Goal: Prevent sender probing for non-existent bandwidth

Problem: “limited to WHAT?”

Added reference to “Right-edge Recovery”

Inject new data when sender receives duplicate ACKs

Goal: Trigger Fast Retransmit/Fast Recovery if possible

Avoid RTO by injecting one or two ACK-clocked segments

Not standards-track - research community feedback?

SLOW: Significant Changes in 02

Recommend against use of TCP timestamp options

Constantly-changing header option prevents header compression

Header compression that compresses SOME options?

Recommend “TCP Buffer Autotuning”

Significant win for hosts with multiple interfaces (LAN + dialup)

Still recommending small receive windows for “ever-slow” hosts

Recommend (at most) 100-millisecond MTUs

200-millisecond delays are human-perceptible (RFC 1144)

200-millisecond MTUs turn off delayed ACKs

ERROR: Significant Changes in 01

Defer LINK topics to LINK

Still working out exact contents of each PILC recommendation!

Restrict scope to TCP

Too many variables for one-size-fits-all UDP recommendation

Remind everyone that errors still hurt TCP performance

No miracle occurs!

Distinguish clearly between standards-track and non-standards-track recommendations

ERROR: Significant Changes in 02

Complete rewrite: Why TCP windows stay small

Split Explicit Congestion and Corruption Notification material

Corruption notification topic swamped by congestion notification

Updated to include Appropriate Byte Counting recommendation

Research/experimentation continues in this area

Added “HTTP and the Dark Side of the Force” section

The WWW has lots of small objects

HTTP/1.0: closed TCP connections a lot

HTTP/1.1: persistent connections intended to be default

This isn't happening, for several reasons

TCPs really won't know what network characteristics are, as long
as connections keep closing

Spencer-the-co-author thinks:

Both documents are scheduled for publication as BCPs after this IETF (per our charter)

At least one more editorial pass is needed for each document (think “03”) before this happens

I’m interested if you have comments

I’m very interested if you see topics that are missing

I’m EXTREMELY interested if you see errors

Your thoughts?